June 5, 2002

CDF Proposals for the Run IIb Project

Introduction

We present here the projects for which the CDF collaboration requests Stage 1 approval. These projects are all aimed at the Run IIb operation of the detector, generally defined as the period when the instantaneous luminosity is expected to reach 5×10^{32} cm^{-2} s^{-1}, and the integrated luminosity is expected to reach 15 fb^{-1}.

The projects needed for Run IIb can be broadly separated into two classes: projects needed due to the increase in instantaneous luminosity, and those needed due to the integrated luminosity and therefore the length of the run required. The former represent an increase in the capability of the detector, beyond the specifications set for the Run IIa upgrade of CDF. The latter are essentially the maintenance costs due to a long run. The Run IIb projects are presented here in three categories: Silicon Detector Replacement, Calorimeter Upgrades, and Data Acquisition and Trigger Upgrades.

Higgs Searches in Run II

The most sensitive channels for the Higgs search (SM or SM-like) in Run II at the Tevatron are the modes $p\bar{p} \rightarrow WH \rightarrow \ell \nu b\bar{b}$ and $ZH \rightarrow \bar{\nu}\nu\, b\bar{b}$. There is small additional sensitivity brought by $ZH \rightarrow \ell^+\ell^- b\bar{b}$ which can be safely neglected for the purpose of this note. These final states will demand two tagged b-quark jets to suppress backgrounds. The main irreducible backgrounds are from $Wb\bar{b}$ and $Zb\bar{b}$, and (in the $\bar{\nu}\nu\, b\bar{b}$ channel) from QCD $b\bar{b}$ dijet production. The trigger efficiencies after the offline analysis cuts for these backgrounds are assumed to scale linearly with that of the signal. We assume the background kinematics is similar to the signal in our approximation.

The signal will appear as a small excess on a large background; one will extract the signal by fitting the observed $b\bar{b}$ mass spectrum to a combination of signal and background. The rates are such that Gaussian counting statistics scaling laws apply rather accurately. At a Higgs mass near 120 GeV/c^2, the Higgs reach improves by 1.5 GeV for every 10% increase in the integrated luminosity. The required luminosity needed for a Higgs measurement, L_{req}, will be reduced if the detector is improved. We can therefore quantify how our Higgs reach is extended by a specific upgrade by estimating the effective reduction in L_{req}. The following general observations can be made about the required integrated luminosity (L_{req}) to exclude or discover the Higgs:

- $L_{req} \propto 1/\varepsilon_b^2$, where ε_b is the *b*-jet tagging rate per taggable jet.
- $L_{req} \propto 1/\varepsilon_{trig}$, where ε_{trig} is the *marginal* trigger efficiency, that is, the efficiency for those signal events which would otherwise pass final analysis cuts which optimize the signal to background.
- $L_{req} \propto \sigma(m_{b\bar{b}})$, the $b\bar{b}$ mass resolution.

In Run IIb, the instantaneous luminosity goal of 5×10^{32} cm^{-2} s^{-1} represents a major challenge to various CDF detector components. The details of the CDF upgrade for Run IIb can be found in the Technical Design Report. We discuss here the impact of the various upgrades on the SM Higgs search.

Silicon Detector Replacement

Motivation

Radiation damage from Run IIa will render the inner layers of the silicon vertex detector inoperable. The main upgrade for Run IIb, a new inner silicon vertex detector, is needed in order to maintain excellent *b*-jet triggering and reconstruction. For the Higgs search the key design goals for this detector are to meet or exceed a 60% *b*-jet tagging efficiency per taggable jet at high p_T, using secondary vertex reconstruction, and to maintain high efficiency in the silicon vertex trigger (SVT).

That a new detector should be built is not an issue; for the SM Higgs search the question is the impact of a reduction of scope (removing a layer, for example) on the Higgs reach. There are two main ways that such a reduction would hurt the Higgs search: loss of *b*-tag efficiency at the trigger level and offline, and loss of redundancy.

Since the required luminosity to discover or exclude the Higgs is inversely proportional to the square of the *b* tagging efficiency, the impact of descoping is quite large. At the highest luminosities, the performance of the inner layers of the COT will be reduced due to the increased occupancy. This effect already leads to some loss of b-tagging efficiency if silicon stand-alone tracking cannot be used. Dropping a layer from the proposed system would lead to a 4% loss of *b*-tagging efficiency, or an 8% increase in required integrated luminosity. Coupling the loss of one silicon layer with the loss of use of the inner layers of the COT, leads to a 13% loss in b-tagging efficiency, which corresponds to an increase of 25 % in L_{req}. Furthermore, the redundancy in the proposed detector will improve the robustness of the system in a multiple interaction environment.

Progress since November, 2001

The first important goal for CDF is to test a prototype stave in its "final" configuration. At the present this will happen at the beginning of September 2002. All stave parts have been ordered and are on time to achieve this goal. Near term activities towards achieving this prototype stave are:

a. SVX4 submitted April 1 and expected back June 11th.
b. Prototype silicon sensors (both axial and stereo) submitted on February and expected back in the middle of July 2002.
c. Hybrid (4 chips) submitted April 20th and expected back June 28th.
d. Mini-port card submitted April 10th and expected back June 28th.
e. Bus cable submitted May 15th and expected back June 15th.
f. All fixtures for prototype module assembly have been produced and tested.
g. All parts for the prototype stave mechanical assembly have been procured.
h. All fixtures for the mechanical assembly of staves have been produced.

Latest Implementation and Critical Milestones

Milestone	Date
Prototype Stave Available	Sept., 2002
Contingency Stave Available	May, 2003
Testing Full Prototype System Complete	Dec., 2002
Testing Full Contingency Prototype System Complete	Aug., 2003
Preproduction Staves Electrical Available	Nov., 2003
Testing Full Preproduction System Complete	Feb., 2004
Stave Mechanical Production Go-ahead	Mar., 2004
Production Staves Available	Apr., 2004
Stave Production Complete	Feb., 2005
Outer Detector Complete	May, 2005
L0 Preproduction Module Ready	Aug., 2003
L0 Preproduction Electrical Test Complete	Dec., 2003
L0 Production Modules Available	May, 2004
L0 Production Modules Complete	July, 2004
L0 Carbon Fiber Support Complete	June, 2004
Inner Detector Complete	Feb., 2005
SVX IIb Ready for Installation	Aug., 2005

The milestones above follow the detector construction flow for both the outer stave layers (which represents >90% of all detector parts) and for L0. The two efforts are taken in parallel and L0 is not on the critical path. The schedule is built following the natural path of prototypes, pre-production and production (whose milestones are listed above) of staves. Stave construction should be maintained as our critical path. The overall schedule is still driven by the SVX4 chip. The first prototype chip was submitted at the beginning of April. The most recent update on the processing status at TSMC shows the chip ready to be shipped back to Fermilab by June 11^{th}. The schedule conservatively assumes that the chip we will receive June 11^{th} will not work properly (or/and the major stave parts such as the hybrid and the mini-port card). In the schedule there is provision for a second submission of the chip, hybrids and mini-port cards. The second chip submission (a minimum "engineering" run if the first fails or a pre-production run in case of a success) occurs in October 2002 together with the "contingency" submission of hybrids and mini-port cards. There are consequently ~7 months between stave availability for testing in the two possibilities. These months could be in part recovered if the first SVX4 chip submission is successful. Following the April 16^{th} director's review recommendations we explicitly inserted some contingency at the lowest level in the schedule. The present schedule has the SVXIIb detector ready for installation in August 2005.

Estimated Cost

As noted in the excerpt we have simplified the silicon detector design in order to limit schedule risks and cost. We have dropped the 90^{o} stereo sensors and extended the same outer stave concept to the difficult radial position of layer 1. The final design is a six layer device with three small angle stereo (1.2^{o}) layers. At present there is no clear path to significantly reduce the cost of the detector except by reducing the number of layers. Eliminating a layer will significantly reduce our b-tagging efficiency which in turn will reduce our Higgs search sensitivity by 8%. The total cost estimated in April, 2002 is $12,473K, with 44% contingency ($5454K).

Calorimeter Upgrades

Central Preradiator Replacement

Motivation

Attaining the best possible $b\bar{b}$ mass resolution will require sophisticated energy flow algorithms which rely on an array of information sources about particle energy depositions. For a detector with coarse segmentation such as CDF, the best emerging approach is to analyze energy flow tower-by tower, using information from the calorimeters, tracking, muon systems, shower max detector, TOF and preradiator. The Run II Higgs report predictions for the SM Higgs reach assume that a 10% $b\bar{b}$ mass resolution can be attained. This goal is ambitious, and close to the physics limit from jet fragmentation and *b* decay effects.

The Run I (and Run IIa) Central Preradiator (CPR) is based on proportional chambers with wires running half the length of the wedge. At the luminosities anticipated in Run IIb, these chambers will not be able to effectively deliver useful information due to high occupancy stemming from the fact that the electronics integrate over several beam crossings.

The preradiator plays a key role in high- p_T photon and electron identification, and in correcting for energy lost by particles showering in the 1.1-interaction-length magnet coil. Studies have shown that for the electromagnetic fraction (30%) in typical jets, an efficient and pure preradiator signal can improve electromagnetic resolution by 20% and may provide an overall improvement in jet mass resolution of a few percent. This is an important part of the overall 30% improvement over Run I that is needed. Identification of electrons in jets can also aid in identifying and correcting the energy of *b* quark jets, which is crucial for the Higgs search.

Concurrently with the CPR upgrade, the present crack chambers (CCR) would be replaced by scintillator-based detectors. These would add to the array of information useful to correct the $\not{E}_T$ measured by the calorimeter and to correct jet energies for losses in the cracks.

Progress since November, 2001

Since November, 2001 the scintillator choice for the preradiator has been studied and changed from the plans presented at that time. The current plan is to obtain sheets of new scintillator from the Russian supplier CDF used for the muon counters installed for Run IIa. This material will be machined into tiles at the Fermilab scintillator facility located at Lab 8 in the village. Our previous plan had been to use excess material from the MINOS project. The new detector approach will provide a detector whose segmentation is more appropriate to CDF. There should be no significant difference in cost between the two plans, since the free MINOS material forced a much larger inventory of optical fibers, due to the small size of the scintillator. Savings in optical fibers will compensate the added cost for the machined tiles.

Latest Implementation and Critical Milestones

The installation of the CPR requires extraction of the central calorimeter arches, which forces the installation to occur in the assembly hall, at the time of the silicon detector installation. Space constraints preclude the scintillator installation in the collision hall. Consequently, the scintillator packages and clear optical fiber bundles used in the readout must be available at the time of the assembly hall work. The latest date for delivery of these parts is then driven by the silicon schedule. Currently, the silicon detector schedule calls for the central detector to be in the assembly hall beginning August, 2005. Key milestones for the central preradiator are listed in the table below. Both our current target dates and the latest dates estimated from the silicon detector constraints are listed.

Event	Target date	Latest date

Procure parts	March, 2003	December, 2003
Prepare scintillator tiles	October, 2003	July, 2003
Assemble modules	March, 2004	December, 2004
Complete fiber bundles	June, 2004	April, 2005

Estimated Cost

At the time of the Director's review in April, 2002, the total estimated cost of the Central Preradiator Replacement was \$806K with 34% contingency (\$273K), of which \$373K will be covered by DOE funds. The Japanese contribution to this project is estimated to be \$273K, and the Italian contribution is expected to be \$160K.

Electromagnetic Calorimeter Timing

Motivation

An additional upgrade to the calorimeter system that is desired for Run IIb is to instrument the electromagnetic calorimeter with timing, similar to the capability already in place for the hadronic calorimeters. Experience with the Run I data indicated that timing information is a useful technique for reducing backgrounds due to cosmic rays, which can be significant in searches for exotic final states that contain photons and missing transverse energy. We have examined the effect of adding electromagnetic calorimeter timing in an analysis searching for associated gravitino pairs and photon production in events containing a photon and large missing transverse energy. We estimate that the limit on the supersymmetry breaking scale $\sqrt{F}$ would be extended from 241 GeV to 260 GeV in Run IIb after adding electromagnetic calorimeter timing.

Progress since November, 2001

Since the November, 2001 meeting of the PAC a new approach towards the implementation of timing in the electromagnetic calorimeter has been pursued. The plan at the time of that meeting, and described in the CDF Technical Design Report for Run IIb is based on modifying the central calorimeter bases, to extract a dynode signal for timing. We have abandoned this approach, and have adopted a strategy that involves the use of an inductive signal pickup off the existing anode output. This approach will make the installation much quicker than the base modification approach. More importantly, the risk to the detector inherent in modifying the existing bases will be removed. The risk of the new cabling approach is minimal. Furthermore, it is completely reversible if some performance problem is uncovered. Tests on the new timing system will be performed during detector access periods in the near future. If this system is adopted, the total cost of the electromagnetic timing project will be significantly reduced.

Latest Implementation and Critical Milestones

The cable installation between the first floor counting room and the central detector must occur during the assembly hall installation of the silicon detector. Consequently, the electromagnetic timing cables must be available well in advance of August, 2005.

Estimated Cost

The cost without contingency for the electromagnetic calorimeter timing hardware was estimated at $208K with 50% contingency ($104K) in April, 2002. A contribution of $130K from Italy is expected. No cost estimates have been made with the new inductive pickup cable scheme, but the installation estimate of $41K should be reduced to half that amount.

Data Acquisition and Trigger Upgrades

Modifications to the Data Acquisition and Trigger systems are proposed both to maintain the existing capability of the systems, and to upgrade to higher bandwidths for Run IIb. Replacement of the processors used in the online system, the level 2 trigger and the level 3 trigger are anticipated just to maintain the existing system. These expenses of operation will be listed first. Upgrades to the systems we believe are needed for the Run IIb instantaneous luminosity include an improvement to our track based trigger, the XFT, an upgrade to the event builder switch and a replacement for the time-to-digital converters (TDC's) used in the drift chamber (COT).

Online Data Acquisition Computing for Run IIb

Software Updates

We depend on several non-Fermilab products: SmartSockets, Java, Linux, VxWorks, and Oracle. We require software and hardware support for some of these critical items.

Computing Infrastructure

The online system currently maintains approximately 50 Linux PC's, 7 mid range file servers for code development and data logging and a several Windows NT based PC's used for iFix. We anticipate having to replace aging systems as needed and it is likely that parts of this system as well as the B0 network will have to be updated for Run IIb. We also plan to replace obsolete front end processors.

We will likely have to replace the network switch and have already had to replace failed switches for Level 3. For Run IIb we will want to replace the SGI file server. This server currently holds the user disks as well as critical online software.

There are currently ongoing discussions about moving the database from Sun to Linux. This would require purchasing three servers and sufficient disk for the database.

Estimated Cost

Software License / Support	
Oracle	$35K/yr
Hardware Support	$5K/yr
Software Support	$5K/yr
Total	$45K/yr

The following is an estimate of the hardware costs for Run IIb.

Maintenance Cost	
20 PCs	$50K
File Server	$100K
Networking	$100K
Front End Processors	$50K
Database	$120K
Misc.	$50K
Total	$470K

Level 2 Trigger Decision Crate

The CDF L2 trigger consists of about 30 VME crates of electronics. It is designed to receive an input of 40-45 kHz and pass on to Level 3 a maximum rate of 300 Hz. At present, the system has achieved about 6 kHz with a physics table that is very close to final form. The entire L2 system should without any modification achieve 20 kHz based on several measurements performed by the Level 2 group. There is a lot of work to do to achieve the final Run IIa performance, but there are no known hurdles at this point that could prevent us from reaching this goal. A lot of progress has been made in the Level 2 trigger in the last two years.

The Level 2 trigger must perform well for Run IIb, as has been stated many times. The Higgs triggers consume about 600 Hz of bandwidth and this does not account for increased fakes in the tracking system or SVT (secondary vertex trigger) trigger system. Fully 80% of our Higgs triggers require a displaced track, and fakes are of great concern to us because we are only beginning to be able to measure them in Run IIa.

The L2 decision crate contains six different custom interface boards (maintained by three institutions) and four alpha processors. Thus far, all the interface boards are working, with the exception of the muon interface board, which we hope will be ready later this year. We are currently running with one alpha processor because getting the full complement working has been a large task for an already overextended L2 trigger group. Running with just one alpha processor has not been a limitation thus far.

The alpha is a 500 MHz processor with a 33 MHz PCI bus, originally designed and built by DEC in the middle 1990's. We are running without an operating system and use the DEC supplied native alpha software. This works fine but the long maintainability is of concern to us. The alpha processors are housed on custom printed circuit boards designed at Michigan according to DEC layout rules. These processors communicate through two busses, the CDF specific VME backplane and the Michigan designed full custom "Magic Bus" that interfaces to the alpha processor PCI bus. The magic bus (backplane in the L2 decision crate) has been difficult to commission because of the design and the high data rates sent over the backplane. We have recently redesigned and built a new magic bus backplane and changed protocol from PECL to TTL in order to reduce the cross talk on the arbitration lines. This has improved the performance although we have made modifications to all the interface boards as well as the magic bus.

It is evident to us that these processors will need to be replaced due to the lack of our ability to build and maintain enough spares (parts are already unavailable on these boards) for a physics program that will end in 2009. We are at present considering two options for replacement. The first is using the D0 beta board. This will require some modifications in order for these boards to work for CDF. We have been in contact with D0 on these points. If we plan to go in this direction, we will need to get their Cadence specific schematics and layout design files. We estimate about two months to understand and modify these files. The beta is an interface between VME, magic bus and a commercial processor. We would then have to fabricate a small number (10) of these boards and put together a group that would perform the debugging, installation, and maintenance of the system.

Another option being considered uses our PULSAR "universal" test station board presently in the final stages of design. Some auxiliary daughter boards are being tested. The PULSAR would allow us to replace all six custom interface boards and bypass the magic bus, providing significant simplification to the present system. The data transfer from the PULSAR to a commercial processor uses the CERN designed (available now from CERN stores) S-link interface. A PULSAR group exists and they are preparing a proposal to define this approach. This system, because it uses a universal board, which is easily reconfigurable, is easy to maintain. Presently we take date with two level 2 decision crates, where one is a “hot spare”. Either option can be implemented with minimal impact on data taking.

Estimated Cost

Either approach is not very expensive. We have estimated about $150K. At this time, a 100% contingency on this project is appropriate.

Level-3 Trigger PC Farm Upgrade

The conceptual design of the Level-3 Trigger PC farm will not need an upgrade since it is can be expanded by adding additional processor nodes. Estimating the real CPU power needs of the PC farm in Run IIb is quite uncertain. We believe it will be mandatory to replace PCs after a three year period since maintenance on older computers is rarely cost effective, compared to replacement. We assume that the number of PCs will be kept constant. Maintenance replacement will then result in a natural growth of the total Level-3 CPU power, since the CPU power per computer increases by roughly a factor of three every two years. The Level-3 PC Farm Upgrade is therefore maintenance of the existing system.

Estimated Cost

The cost involved in the replacement of PCs is $160K per year including $30K contingency. We anticipate four years of replacement (2002-05 inclusive).

Higgs Trigger Strategies and Bandwidth Upgrades

CDF plans to exploit both calorimetry and displaced track based triggers to optimize our sensitivity to Higgs searches. The strategy is to use a calorimeter-based $\not{E}_T$ trigger for the process $ZH \rightarrow \bar{\nu}\nu\, b\bar{b}$, and use the SVT-based trigger with relaxed calorimeter cuts to control the trigger rates needed for the background studies of $b\bar{b}$ dijet events and to complement the calorimetry based trigger. Table 1 shows the minimal set of triggers designed for Higgs search at Run IIb and their estimated level 2 trigger rate at 5×10^{32}cm^{-2} s^{-1}, based on linear extrapolations of the Run I measured rates. The total trigger rate shown below exceeds the present 300 Hz bandwidth limit.

Triggers	*Rate (Hz) at $5\times10^{32}cm^{-2}\ s^{-1}$*
Calibrations	
Minbias	12.5
Jets	43.5
Lepton+SVT	45.5
$J/\psi \rightarrow \mu^+\mu^-$	37.5
$Z \rightarrow b\bar{b}$	16.0
Higgs	
High p_T CEM+PEM	92.5

High p_T CMX+CMUP	100
$\not{E}_T$ + 2 Jets	55
Multi-Jet + SVT	47.5
$\tau + \not{E}_T$, τ	45.5
High p_T Jets +SVT	100
SUSY, etc.	
High p_T Isolated Photon	50
Ultra high p_T Photon	20
High E_T di-photon	10
Total	**676**

With nearly five interactions per crossing at a crossing rate of 132 ns, high detector hit occupancies lead to fake trigger objects, leading in turn to nonlinear growth in trigger rates. Run I experience showed that the onset of nonlinearity was difficult to predict and oftentimes a large effect. The implication of this is that to be assured of maintaining Higgs sensitivity, and to allow for unforeseen limitations in the trigger, we need to build into the DAQ system sufficient bandwidth to handle these situations. The proposed design would attain 1 kHz, which provides some margin to accommodate nonlinear rate growth.

Upgrade of the Event Builder Switch

Motivation

The maximum bandwidth theoretically achievable with the existing Run IIa system is 240 MBytes/s. In practice about 60% of this limit has been achieved in benchmarking tests with simulated data sizes corresponding to the expected detector occupancy. It is possible that the performance under test conditions can be further improved to about 80% of the theoretical limit after tuning the system. The performance when processing real data on the other hand depends crucially on the load balancing among the various ATM switch inputs and it is desirable to assume that at best 50% of the theoretical bandwidth can be used.

To match the requirements quoted in the introduction of this section corresponding to a sustained data rate of at least 250 MBytes/s the existing system is not sufficient and an upgrade is necessary. For the upgrade a sustained bandwidth of at least 400 MBytes/s is required to take into account load imbalance in the ATM inputs and fluctuations in the data size.

A simple straightforward upgrade scenario is therefore proposed, in which the same technology is used. The existing ATM switch uses OC3 connections, which are upgraded to the four times more powerful OC12 connections. This also implies a complete replacement of the switch. The main work will be to rewrite the low level drivers, which are not exactly suited for our application.

Estimated Cost

The cost was been evaluated in April, 2002 and estimated as \$510K with 35% contingency (\$178K).

Addition of Stereo Data to the XFT

Motivation

The CDF Level-1 track trigger (XFT), which can be applied to ~80% of all Run IIb triggers for high p_T physics, suffers from a rapid increase in fakes once the number of overlapping minimum bias events exceeds something like 5-6 interactions per crossing. CDF has applied the full Level-1 track trigger hardware simulation to a sample of $t\bar{t}$ Monte Carlo Events. Minbias events are overlapped and the fake rate is determined. It is found that the fake rate increases substantially as a function of the number of interactions per crossing. The results are shown below. The impact of an increased number of fakes affects directly the number of fake single electron and muon triggers, which combine to use 25% of the trigger bandwidth in Run IIb.

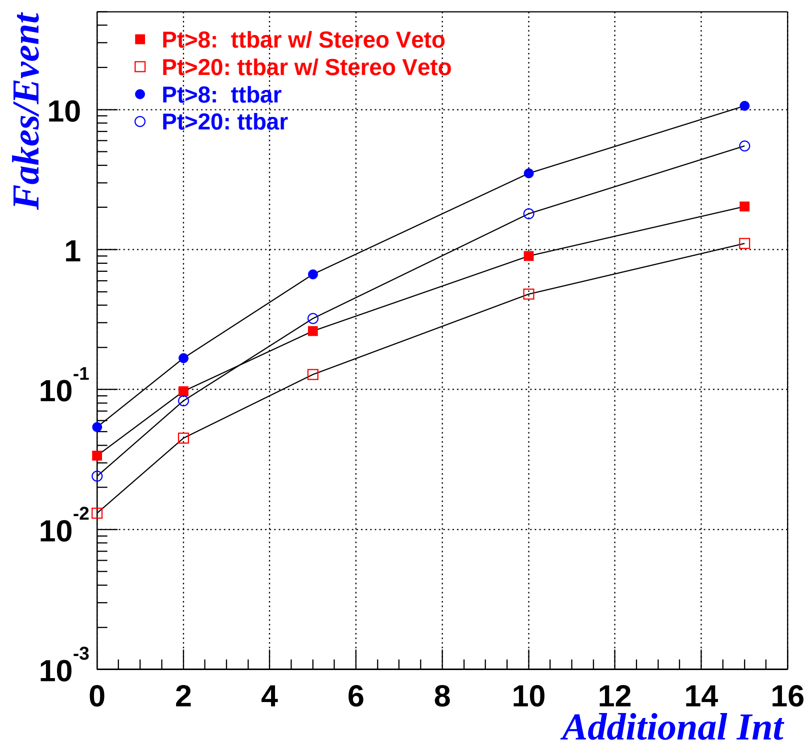

Figure 1. The average number of fake tracks per event in ttbar Monte Carlo as a function of the number of additional interactions per crossing. The curves are shown for XFT tracks with Pt>8 GeV/c and Pt>20 GeV/c. Also shown is the reduction of fake rate provided by requiring the

presence of a segment in the outer stereo layer of the COT. The fake rate at large number of interaction (N>8) is almost entirely driven by the soft additional interactions, not the hard scatter (in this case ttbar).

The SVT relies on the XFT to provide track roads from the COT to the silicon vertex detector. In Run IIb the high rate and occupancy lead to an increase in fake tracks, and eventually, with more than about five interactions per crossing, the angular resolution of the XFT degrades to the point where the SVT performance is compromised. Most of these fakes are due to random crossing of low momentum tracks.

Extrapolating from current experience, we expect that the SVT trigger efficiency will reach ~90% in the good silicon regions. A similar performance should be achievable with the new vertex detector, however, at high luminosity the poorer XFT performance (resolution and fake rates) could degrade this significantly.

To limit the rate of fakes, we propose to instrument the outer stereo layer of the COT with XFT electronics. Segments found in the stereo layer will be matched to those obtained in the r-ϕ view. The initial studies indicate that the upgraded XFT stereo information can reduce the fake rate by a factor of 2.5 at a luminosity of 5×10^{32} cm^{-2} s^{-1}. As Figure 1 shows, this could extend the maximum useful instantaneous luminosity by a factor of two, which translates directly to reduced running time.

The simulation which was used to predict the luminosity dependence of the XFT fake rate does not include the observed factor of two increase in the occupancy of the COT, compared to expectations. As a consequence, we are concerned that the increase in the fake rate may occur at a much lower number of interactions per crossing than indicated above (approximately 5). This could lead to a effective increase of L_{req} proportional to that fake rate – perhaps 20%.

Estimated Cost

The cost was been estimated as $400K with 50% contingency ($200K).

TDC Replacement

Motivation

The Central Outer Tracker (COT) uses a custom 96 channel TDC designed and built by the University of Michigan. The COT uses 315 TDC's not including spares. After considerable problems in early versions of the board, the system has been very stable in the last year and CDF is very satisfied with its performance. We do however have a shortage of boards, and are awaiting an order for 100 additional boards to be placed this summer by the Michigan group. There are at present only seven spares in the system. In the meantime, repairs to faulty or failed boards are taking place at Michigan. The repair effort should provide CDF with about 20 more boards, not enough to satisfy the needs of the COT, which presently operates with about 40 single channel problems in the system. The muon systems in CDF employ about 100 TDCs. Again, we are lacking in spares but the system performs well. Once the new boards arrive, the TDC system will meet our needs for Run IIa.

CDF has recently reviewed the TDC system and has found that the readout speed is not adequate for our 1 kHz readout specification for Run IIb. The reason for this is that the COT occupancy is higher by a factor of two in minimum bias events and a factor of three in "physics triggers" events (CDF Note 5889). The readout speed of the TDC system is determined by the slowest TDC in a given VME crate. The inner radius COT layers therefore take longer to readout than the lower occupancy outer COT layers. An internal study (CDF Note 5824) of the readout of the DAQ system has found that the CDF detector is limited by the readout speed of the TDCs in the COT for Run IIb. All other systems in CDF can meet the 1 kHz readout rate specification. The recent review of the TDC system (May 23, 2002 K. Pitts Chair) has recommended that CDF pursue several noninvasive options (eg. improved DSP code) for increasing the speed of the present TDC to try and meet Run IIa and Run IIb needs. The committee had significant concerns that the present TDC system would not meet the Run IIb needs and therefore also recommended that CDF pursue the design of a much faster TDC system based on the high speed Altera Startix FPGA. A rough (although elegant) design using FPGAs was presented to the committee by Harold Sanders from the University of Chicago.

Specifically the committee recommended:

1) Continue to improve and gain experience with the Michigan TDCs.
2) Continue with new TDC R&D, including adapting the design for the CDF COT specific application and considering new readout techniques to overcome data rate/volume limitations. A rough cost estimate was presented. To replace all COT TDCs would require about $900 K including M&S, engineering with 100% contingency ($900K).

In summary, CDF is pursuing a new, faster readout TDC system that will allow us to reach the 1 kHz goal for Run IIb. It is felt the present system will meet the Run IIa goal of 300 Hz. The system cannot meet the 1 kHz readout rate for Run IIb due to the increased occupancy in the COT.

Conclusions

The Higgs search needs the integrated luminosity attainable in Run IIb, and a detector capable of exploiting it fully. The effect of the proposed Run IIb upgrades is summarized here:

- A new silicon vertex detector is the essential element of the Run IIb upgrades, and must deliver a *b*-tagging efficiency of at least 60% of taggable high-E_T jets. A loss of even 10% (to 54%) in this efficiency due to lack of redundancy would cost 20% in required integrated luminosity to discover the Higgs, which translates into 3 fb^{-1} for discovering a roughly 120 GeV Higgs, more than a half year of running time.

- The upgraded detector needs the best possible $b\bar{b}$ mass resolution, and this in turn demands maximal information regarding jet energy flow. The proposed CPR and CCR upgrades will enhance the information available in the Run IIa data, and likely lead to a few percent improvement in jet energy corrections overall. This detector improvement is equivalent to several months of running time towards the discovery of the Higgs Boson.

- The timing upgrade to the electromagnetic calorimeters provides an improved limit on the supersymmetry breaking scale $\sqrt{F}$ which would be extended from 241 GeV to 260 GeV in Run IIb.

- Triggers for the isolated lepton and $b\bar{b} + \not{E}_T$ final states must remain fully efficient at high instantaneous luminosities. Instrumenting the outer COT stereo layer for use in the XFT decision would limit fakes and aid in hit pattern recognition in the SVT. We estimate that this upgrade will extend our track trigger effectiveness to the instantaneous luminosities anticipated in Run IIb. Without this upgrade 20% more integrated is required for the Higgs discovery.

- Finally, the upgrade in bandwidth for the Event Builder from 300 Hz in Run IIa to 1 kHz in Run IIb, along with a complementary upgrade to the TDC system, allows us to retain the high efficiency for the triggers on which the Higgs search is based. The proposed 1 kHz system is well-matched to the rates and signal efficiencies of the most important high- p_T physics of Run IIb

Summary of Milestones

Silicon detector upgrade Milestones	Date
Prototype Stave Available	Sept., 2002
Contingency Stave Available	May, 2003
Testing Full Prototype System Complete	Dec., 2002
Testing Full Contingency Prototype System Complete	Aug., 2003
Preproduction Staves Electrical Available	Nov., 2003
Testing Full Preproduction System Complete	Feb., 2004
Stave Mechanical Production Go-ahead	Mar., 2004
Production Staves Available	Apr., 2004
Stave Production Complete	Feb., 2005
Outer Detector Complete	May, 2005
L0 Preproduction Module Ready	Aug., 2003
L0 Preproduction Electrical Test Complete	Dec., 2003
L0 Production Modules Available	May, 2004
L0 Production Modules Complete	July, 2004
L0 Carbon Fiber Support Complete	June, 2004
Inner Detector Complete	Feb., 2005
SVX IIb Ready for Installation	Aug., 2005

Calorimeter – CPR upgrade Milestones		
Event	Target date	Latest date
Procure parts	March, 2003	December, 2003
Prepare scintillator tiles	October, 2003	July, 2003
Assemble modules	March, 2004	December, 2004
Complete fiber bundles	June, 2004	April, 2005

DAQ and Trigger – Stereo XFT upgrade Milestones		
Event	Target date	Latest date
Procure parts	Sept., 2002	Sept., 2003
Install electronics	April, 2003	May, 2004
Install cables	Sept., 2003	Oct., 2004
test	June, 2005	June, 2005

Estimated Costs

Project	Base Cost ($K)	Cont.%	Cont. ($K)	Total ($K)
Silicon	12,473	44	5454	17,927
Calorimeter upgrades				
CPR replacement	806	34	273	1079
EM timing	208	50	104	312
Trigger and DAQ upgrades				
TDC replacement	900	100	900	1,800
Stereo XFT	400	50	200	600
Event Builder	510	35	178	688
L2 decision crate	150	100	150	300
Online DAQ software	135	30	41	176
Online DAQ hardware	470	30	141	611
L3 PC farm	520	25	130	650
Total ($K)	16,572	46	7571	24,143